

Attorney's Docket No.: 042390.P5379 Express Mail No.: EM560651152US

## UNITED STATES PATENT APPLICATION

## **FOR**

# EARLY ACKNOWLEDGMENT OF PRIMARY PACKETS

Inventors:

Jerrold V. Hauck David W. LaFollette

Prepared By:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN 12400 Wilshire Blvd., 7th Floor Los Angeles, California 90025-1026 (310) 207-3800

15

20

25

## **BACKGROUND OF THE INVENTION**

## (1) Field of the Invention

The invention relates to data transfer systems. More specifically the invention relates to bandwidth reclamation on a full duplex bus.

# (2) Related Art

For high speed data transfers over a distance, various high speed serial data transfer mechanisms have been developed. The IEEE Standard for a High Performance Serial Bus, IEEE Std. 1394-1995 published August 30, 1996 (1394-1995 Standard) and its progeny provide a high speed serial protocol which permits implementation of high speed data transfers. The existing progeny includes P1394a Draft Standard for a High Performance Serial Bus (1394a Standard) and P1394b Draft Standard for a High Performance Serial Bus (1394b Standard). Generically, systems implementing 1394-1995, 1394a, 1394b or subsequent revisions and modifications thereof are referred to herein as 1394 systems. In 1394 systems a plurality of nodes are organized into a tree topology. Additionally, all nodes are initially deemed peer to peer but on bus reset, one node assumes root status while the remaining nodes become branches and leaves of the topology. Both 1394-199\$ and 1394a are half-duplex protocols. Thus, data may only flow in one direction at any given time. Typically a primary packet is sent by a source node out all ports. A downstream destination node identifies from a header of the packet that the packet is directed to it. The recipient then receives the packet and after the entire primary packet is received and the acknowledge gap permits reversing the bus, an acknowledgment (ACK) packet is sent to the source. If the packet is not completely received, the



10

destination node still waits until the transmission and ACK gap are complete before sending an ACK packet to the source indicating that the source should send the packet again later because the destination was busy or an error occurred.

Both of these protocols employ a discrete arbitration phase during which all nodes seeking access to the bus send arbitration requests to the root node. The root node then resolves the arbitration and grants the bus to one of the arbitrating nodes. During this discrete arbitration phase no other useful work can be done on the bus.

The 1394b Standard provides a full-duplex serial bus protocol. Busses employing the 1394b Standard may be able to reduce the ACK gap as the bus direction need not be reversed, but they still generally follow the primary packet/ACK scheme discussed above. Bus bandwidth is often a scarce resource. Efforts to improve bus efficiency are ongoing.

# **BRIEF SUMMARY OF THE INVENTION**

An apparatus for improving bus efficiency is disclosed. A transceiver is provided. A state machine is coupled to the transceiver to generate a NAK concurrently with the receipt of a packet when the packet cannot be successfully accepted.

# **BRIEF DESCRIPTION OF THE DRAWINGS**

**Figure 1** is a block diagram of a physical layer of one embodiment of the invention.

Figure 2a and 2b are block diagrams of a transaction between a source and destination node in one embodiment of the invention.

Figure 3 is a block diagram of a tree topology of one embodiment of the invention.

Figure 4 is a flow diagram of source node activity in one embodiment of the invention.

**Figure 5** is a flow diagram of destination node activity in one embodiment of the invention.

10

20

25



A source node sends a primary packet toward a destination node along a full duplex bus. If the destination node identifies that it cannot receive the packet, it sends an acknowledgment (ACK) packet in the opposite direction along the full duplex bus towards the source node. The ACK packet contains a code indicating that the packet could not be successfully accepted and should be retried later. The 1394a Standard defines acknowledge codes which stimulate a retry including ack\_busy\_\*; ack\_tardy; and ack\_\*\_error. As used herein, an ACK packet containing such a code, or any other signal that a data transmission will not be accepted and should be retried is termed a "NAK." Upon receiving the NAK, the source node aborts packet transmission of the primary packet which by virtue of the NAK has been identified as futile. Accordingly, the bandwidth which would have otherwise have been used for the futile packet transmission can be reclaimed and used for some other purpose.

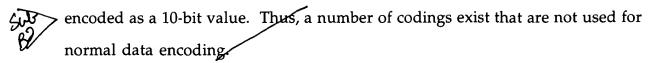
Now referring to Figure 1, a physical layer (PHY) 10 has one or more ports 15 SUD> which can be connected to a serial bus not shown. For each such port, per port logic 12 is provided as part of the PHY/10. The PHY 10 receives a receive signal 32 and transmits a transmit signal 34 through each active port. Per port logic 12 includes a decoder 22 to decode and/descramble the incoming receive signal 32. In one embodiment the decoder 22 is an 8B10B decoder and descrambler, which performs 8-bit 10-bit decoding. This means each 10-bit value coming in as the receive signal 32 is decoded into an 8-bit value. The decoder 22 is coupled to a port state machine 26 which receives a control value decoded by the decoder. Similarly, an encoder 24 is used to encode and scramble the outgoing transmit signal 34. In one embodiment of the invention the encoder performs 8B10B encoding. Analogous to the decode case, each byte of data to be transmitted is

10

15

20

25



Additionally, the decoder 22 is coupled to an elasticity buffer 18. The elasticity buffer 18 buffers data decoded from the received signal 32 while that data is awaiting transfer to the link and/or repetition out on one or more ports of the PHY 10 as the transmit signal 34. The PHY state machine 14 and the arbitration state machine 16 may be combined as a single state machine or may be implemented as two separate state machines. The PHY state machine 14 provides the intelligence for the PHY including generation of any PHY generated packets. The arbitration state machine 16 assesses incoming arbitration information and provides outgoing arbitration requests. If the node is originating a packet on the bus it is the nominal root node. The nominal root node has complete arbitration state information for the topology. Based on this information the arbitration state machine 16 of the nominal root node grants the arbitration request of a highest priority branch. Since nominal root status passes with the grant of the bus, each node should be able to act as arbitrator for the topology. Copending application entitled "Distributed Arbitration on a Full Duplex Bus," Serial No. 09/017,451, describes at least one suitable embodiment of a distributed arbitration system. While much of the subsequent description is devoted to an embodiment in a distributed arbitration topology, the instant invention is also applicable to a non-distributed arbitration topology. For example, even in a system that uses a physical root node to conduct all arbitrations or those systems that continue to employ subaction gaps between each subaction, the invention provides significant bandwidth savings.

Figures 2a and 2b show a sample transaction in one embodiment of the invention. Source node 110 is transmitting a PACKETA 116 to destination node 112. Destination node 112 has determined that it cannot accept PACKETA 116.

Sul

5

10

15

20

25

The inability to accept might be caused by e.g., insufficient available resources. This can often be identified from the packet header, e.g., the packet header indicates that the packet is of a size that would exceed the destination buffer resources. While PACKETA 116 is still being transmitted to destination node 112, destination node 112 transmits a NAK 114 upstream to source node 110. Upon sending the NAK 114, the destination node asserts its arbitration request 120 on the upstream line. In **Figure 2b**, source node 110 having received NAK 114 previously, aborts PACKETA 116 and issues a grant 122 to the only requesting node, here destination node 112. By aborting the packet that must be resent later the remaining packet time may be reclaimed and used for useful work.

In an alternative embodiment, which might be used where distributed arbitration is not available, the source node upon receiving the NAK aborts the packet. Then after a subaction gap all nodes may begin arbitrating through the physical root node. Because the start of the subaction gap occurs earlier in time e.g., bus bandwidth is not wasted continuing to send a futile packet, bus efficiency is improved.

Figure 3 shows a system that includes a plurality of nodes 50-58. The nodes 50-58 are organized in a tree topology. In this topology node 50 is the physical root node. However, employing a distributed arbitration technique, which node is the nominal root, nominal branches and nominal leaves depends on which node is permitted to originate a packet on the bus. The nodes are all interconnected by a full duplex serial bus 100. In one embodiment, serial bus 100 follows the 1394b protocol. Each node has one or more ports 30 through which communication between serial bus 100 and PHY 10 is conducted. The PHY 10 communicates with a link layer 36, which in turn communicates with a transaction layer 40.

10

15

20

25





In Figure 3, node 54 is transmitting a primary packet of a subaction (PACKETA) out all of its ports. Node 54 is therefore the nominal root node and receives arbitration information from each child node such that node 54 has the complete arbitration state of the topology available to it. Because node 54 is transmitting a packet on the bus, the opposite directional path towards node 54 from nodes 50, 56 & 58 is available. In a distributed arbitration case, the other node will forward arbitration state information in this upstream direction. In Figure 3, it is presumed that the destination node is node 50. Node 50 having begun to receive PACKETA in a previous unit of time and identified an inability to successfully accept PACKETA transmits a NAK to node 54. The NAK may be followed by arbitration state information. Upon receiving the NAK in a next unit in time, node 54 will abort PACKETA and grant the bus to another requester. Inasmuch as current primary packet time may approach 42 µsec at 800 Mbps, aborting a packet after e.g. 10 µsec can often result in substantial opportunity to reclaim bandwidth. This is true even where the bus employs a discrete arbitration phase.

Figure 4 shows a flow chart of activity in the source node in one embodiment of the invention. At functional block 150, the source node begins transmitting a primary packet. If at decision block 152 a NAK has been received, the source node aborts the transmission at functional block 154. If, however, no NAK has been received at decision block 152, a determination is made at decision block 156 if the transmission is complete. If the transmission is complete (e.g., an acknowledgment received) or after aborting transmission, the source node grants the highest priority of the arbitration request at functional block 158. If the transmission is incomplete, the source node continues to transmit and watch for an incoming NAK.

10

15





Figure 5 shows a flow diagram of activity in a destination node of one embodiment of the invention. At functional block 170, the destination node begins receiving the primary packet transmitted by the source node. If an inability to successfully accept the packet is detected by the destination node at decision block 172, a NAK is sent at functional block 174. Otherwise, a determination is made if the whole packet has been accepted at decision block 176. If the whole packet has been accepted, an ACK packet is sent at functional block 178 indicating the packet has been successfully accepted. If, however, the whole packet has not yet been accepted at decision block 176, the destination node continues to receive and watch for an inability to complete packet acceptance.

In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes can be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. Therefore, the scope of the invention should be limited only by the appended claims.